Visual reactive collision avoidance for unmanned surface vehicles

Daniel Donavanik, NREIP
ddonavan@cis.upenn.edu
Mentor: Mike Bruch
17 August 2005

Problem

- Existing USV radar navigation is ineffective at avoiding transient and moving obstacles in close proximity (~75m)
- *Goal*: An auxiliary visual system to complement radar navigation
- *Reactive collision avoidance*

Background

Perceptor with stereo pair

- Vision for robots in uncontrolled environments typically uses *calibrated stereo* (CCD) or *range* imaging (commonly laser range scanner)
- Purpose is to easily differentiate and assign depth information to objects in a scene based on visual data
- Both of these require precision optics and static calibration for extraction of 3D geometry

Approach

- A more dynamic, flexible system to obviate the need for static calibration
- A system which operates entirely in image space, with minimal reliance on physical parameters
- A monocular approach
  - No inherent way to extract depth information from a single 2D image, *but* can approximate using relative distances within the frame
  - USV specific: **horizon finding**

Approach

- Obstacle positions may be triangulated given their pixel distances from the *horizon*: the interface between water and sky/shore.
- All calculations are done in image space.
- **Two tasks:**
  - Calculation of a *parameterized horizon*
  - Extraction of obstacles on water surface

Approach (concept)

- Triangulation performed using distance from horizon, offset from center.

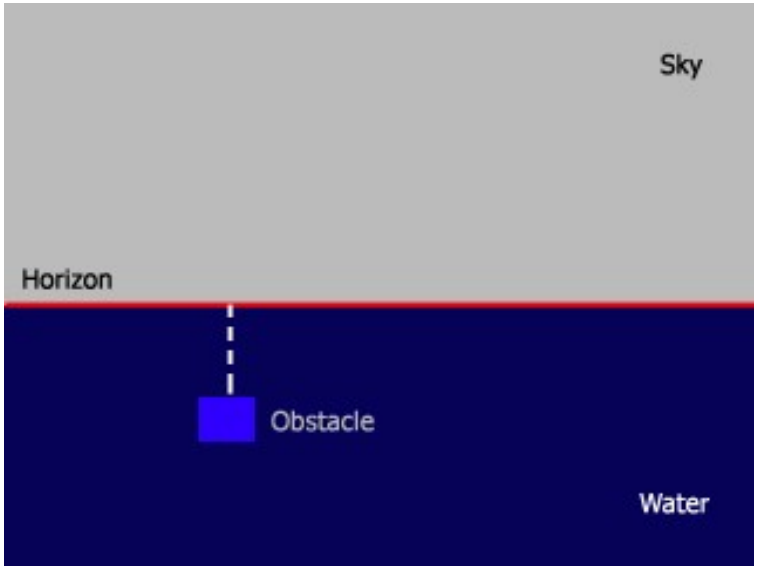

USV view

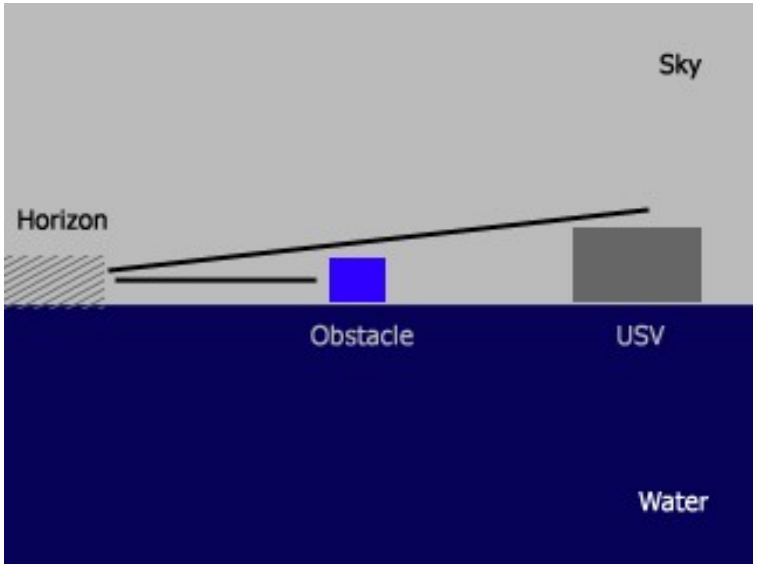

Side view

Specification

- System comprises two modules working alongside the present USV systems
- *Inputs*
  - Video sequence
  - Orientation as provided by boat sensors
  - Distance to horizon (from nautical charts)
- *Output*
  - *3D location of obstacles (in some appropriate coordinate system), to be sent to path planner*

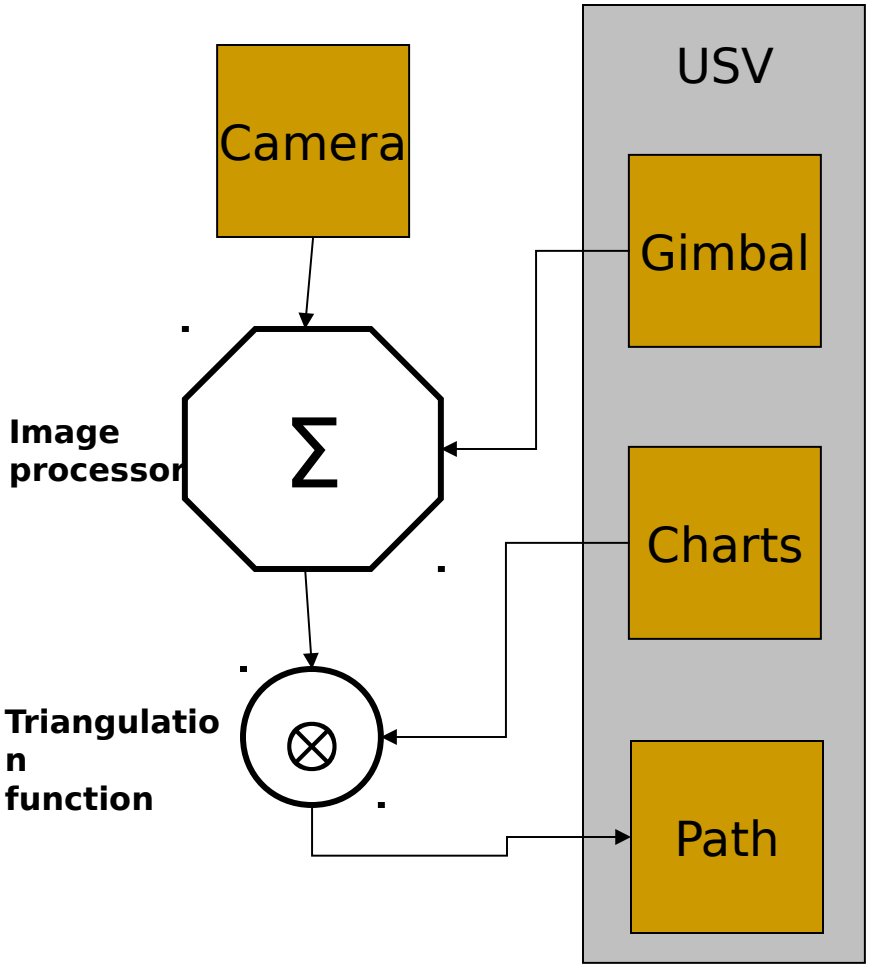

Challenges

- Accuracy
  - Optics must be relatively free of distortion
  - Thermal vs. light intensity imaging?
- Speed
  - Image processing and triangulation subsystems must be able to process images in real time (at a reasonable frame rate)
- Integration
  - The system must be portable enough to be incorporated with the existing USV infrastructure

Module I: Horizon segmentation

- A very recent area in computer vision
  - Nearly all work focuses on UAV applications
- Extraordinarily difficult
  - For UAV, only need an approximate solution giving cues to orientation
  - USV requires accuracy enough for pixel-based measurements!

Possible approaches

- Gradient based techniques
  - Averaging causes problems, inaccurate (see right)
  - Many gradients along any vertical column

M.C. Nechyba et al

- Assumption of discrete classes
  - *Cornall, Egan et al (2005):* Use class centroids
  - Radial system requiring “fisheye” optics

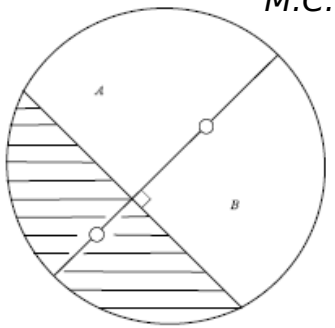

Figure 1: The horizon is perpendicular to the line joining the sky and ground centroids

Cornall & Egan

Linear discriminant analysis (LDA)

- General class of algorithms which tries to find the *discriminant* giving rise to the optimal 2-part partition
  - *Todorovic & Nechyba* (2003): Multiscale linear discriminant analysis
    - Aesthetically similar to Warnock's graphics algorithm
    - *Computationally expensive*
  - Ettinger et al (2001, University of Florida):
    - UAV navigation
    - Assume two-part partition of image between sky and ground
    - Iterate over prescribed range of discriminant parameters
    - Minimize *RGB color variance* for each class

Linear discriminant analysis (LDA)

- Concept: algorithm “scans” over range of possible arbitrary discriminants
- Range is over elevation (pitch) and slope (roll).

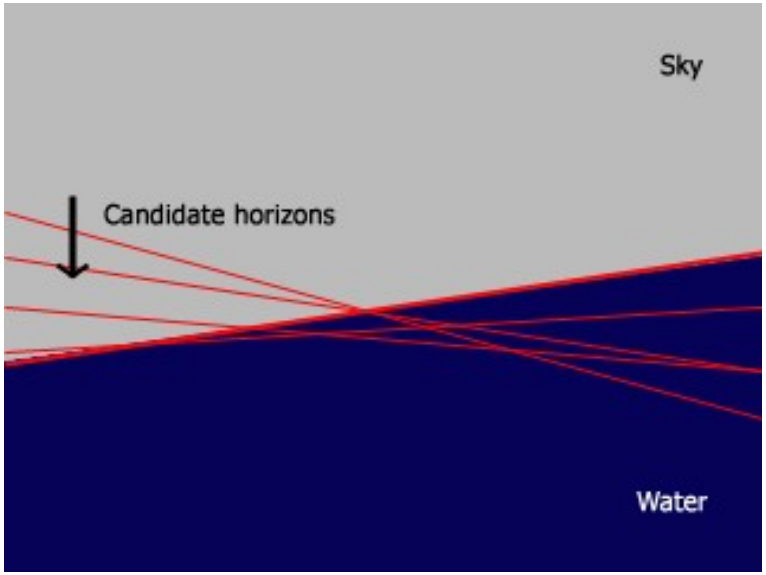

Linear discriminant analysis (LDA)

- Ettinger algorithm
  - An optimization problem:
    - Maximize the inverse sum of the determinants of RGB covariance matrices for both classes as given by arbitrary partition
  - Assumes consistency within “water” and “other” classes
  - Not concerned with gradients or other features
  - Computationally expensive, *but* highly scalable
    - Running time proportional to (1) image resolution, (2) ranges of pitch, roll* parameters (2) incremental step
    - Speed and accuracy are neatly traded

Horizon results (Ettinger al

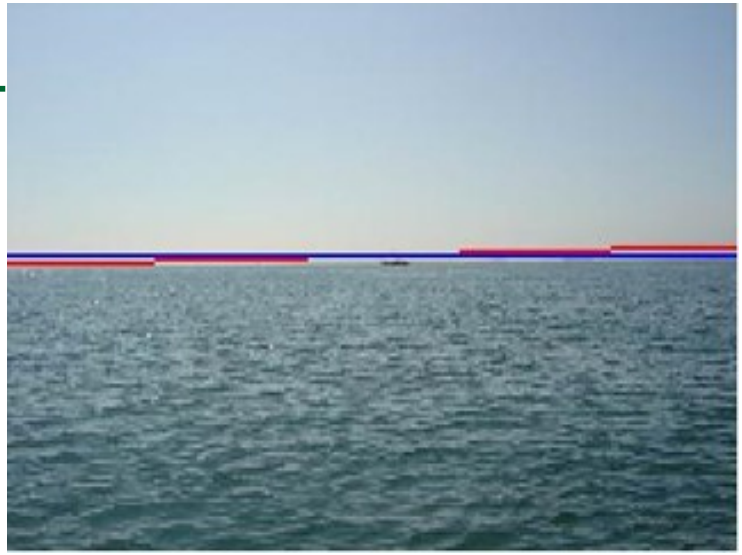

Problems

- Extremely slow
  - Loss of detail with reduced-resolution images
  - Slow anyway: 30 seconds per frame! (*MATLAB* implementation)
- Not accurate
  - Fine for UAV (Needs spatial orientation only), but not for pixel measurements
  - Gets confused by image features in multiple regions
    - Typically, “sky” class extraordinarily heterogeneous, may itself be partitioned

Revised approach

- Ettinger algorithm ignores useful information given by *lines* and *features* present in the image

- **Solution**: Use line-finding in tandem with Ettinger partition optimization
  - Fast *Hough transform*
  - Reduced search space

Horizon results (Revised al

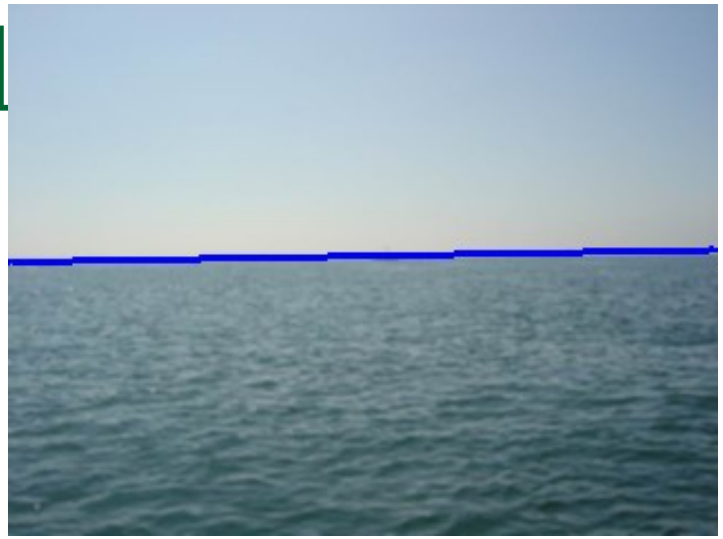

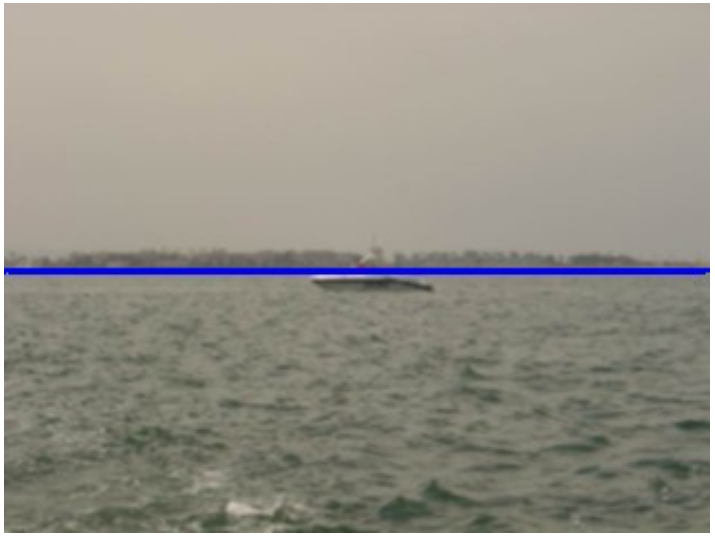

And very fast, too!! (<<1 second, at full resolution!)

Horizon results (video)

- Satisfactory real-time performance
  - *C++ implementation*
  - *Pentium IV, 512MB RAM*
  - *320 x 240, 15fps, Cinepak codec*

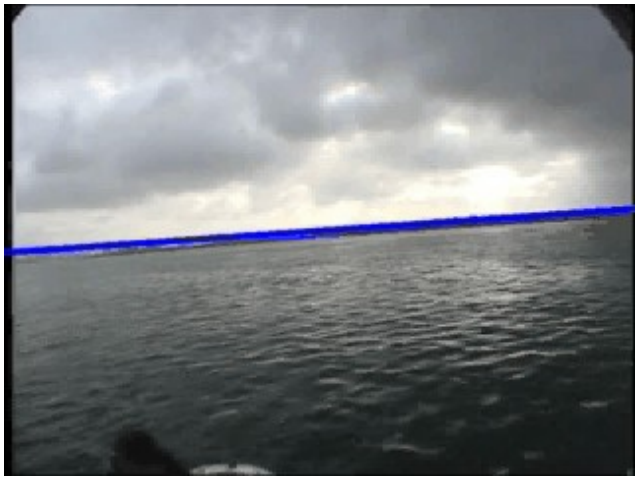

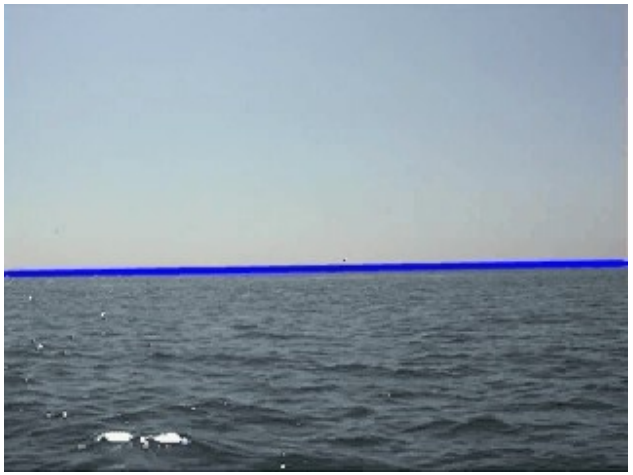

Final remarks

- Optimization function designed to take advantage of RGB color information, *but can be modified for use with grayscale images*
  - Use max eigenvalues of covariance matrices

$$J = \frac{1}{|\Sigma_s| + |\Sigma_g| + (\lambda_1^s + \lambda_2^s + \lambda_3^s)^2 + (\lambda_1^g + \lambda_2^g + \lambda_3^g)^2}$$

- Modified optimization gives priority to most homogeneous class
  - Increase multiplicity of either “sky” or “water” terms
  - Can toggle at run time

Module II: Obstacle segmentation

- Obstacle segmentation on water surface
  - Differentiation among discrete objects on the water surface
  - Differentiation between true obstacles and transient artifacts (glare, froth, insignificant solids)

- A very hard problem
  - Not like traditional image segmentation
  - Water is a time-varying *dynamic texture*
  - Time-tested techniques (background subtraction, etc.) don’t work

Approach

- Mathematical techniques for representing/ dealing with dynamic texture
  - *Statistical estimation (Markov processes, etc.)*
  - *Kalman filtering* (Zhong et al, 2003)

Zhong et al

- Derive a *generative* (rather than static) model of the background based on a training set of background images
  - Requires training (obviously)
  - Unsupervised learning not possible

Approach

- The need for *supervised learning* undermines the flexibility/spontaneity of the one-camera approach
  - Surface may have very different appearance depending on location, time of day, weather
  - Representative training images for all conditions?

- **So**: Avoid using explicit texture model

Optical flow object detection

- *Snyder et al* (2004): water surfaces provide '*sparse, short-lived and non-rigidly moving flow.*'
  - Standard corner detection reveals surface features
  - Tracking features over a set interval reveals coherent objects (transient features disappear)

- Pyramidal Lucas-Kanade algorithm
  - Fast, sparse feature-tracking algorithm
  - Very customizable
  - Features exhibiting coherence in time, space (proximity) and motion (vector direction / magnitude) may be clustered

Optical flow object detection

- Where *I(t)* is the frame captured at time *t*:
  - Detect corners in $I(t_0)$ below the horizon
  - From $t=t_0..t_n$, n ~10-30, track features using iterative application of Pyramidal Lucas-Kanade algorithm
  - At time t_n, cluster remaining features into regions using standard region growing based on proximity and direction
  - Region centroids taken as obstacle centers

Obstacle detection results

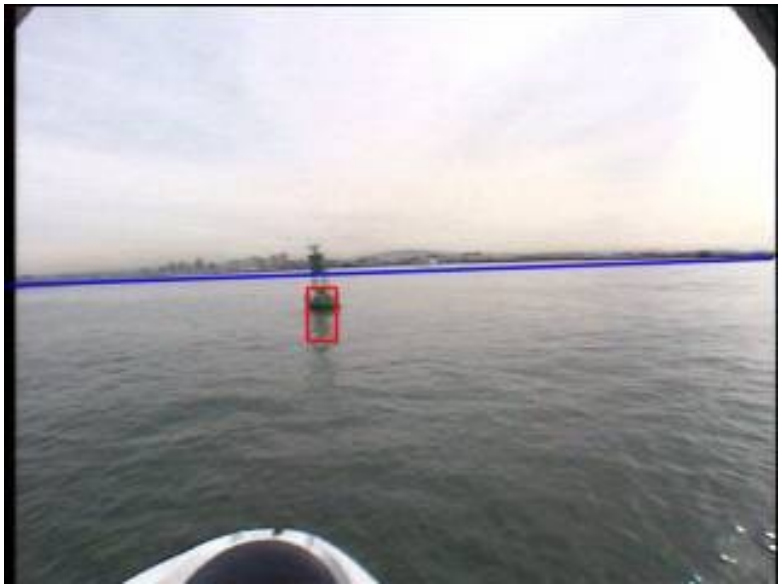

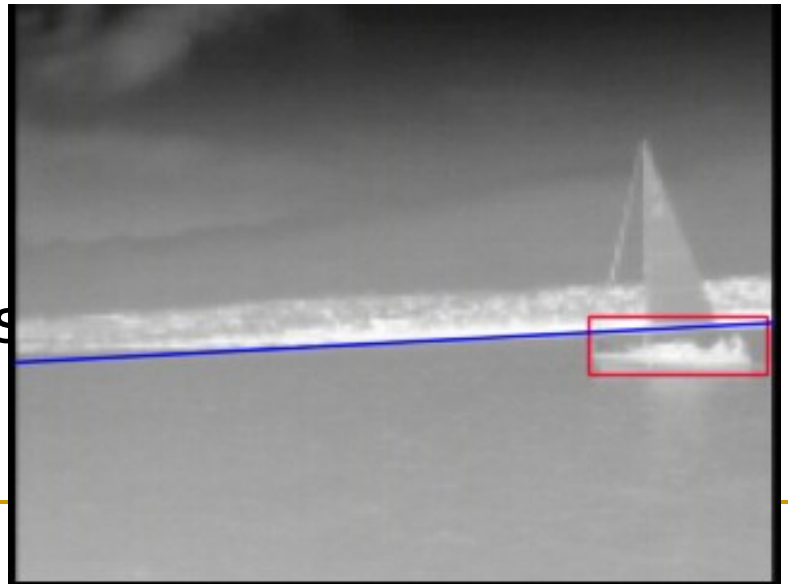

- Still very fast
- Also: *Thermal imaging* (right) resolves ambiguities

Obstacle detection results (video)

- Red boxes indicate “hits” for a given region of the frame
- Note effect of radial lens distortion
- Thermal imaging detects moving objects well, gets confused by horizon
  - Hybrid system suggested

Module III: Navigation output

- Region coherence
  - Regions which consistently appear over an interval represent threatening obstacles
- Output to navigation
  - Either treat obstacles individually (hard), or divide the frame into periods:
    - Identified obstacles contribute to the “score” for a region of the frame in a weighted average
    - High-scoring regions represent a high probability of occlusion.

Further work

- “Intelligent” configuration
  - Highly customizable system with parameters which may vary optimally given prior knowledge about prevailing weather conditions, boat orientation, location, etc.
  - Feature detection, line finding, horizon detection, and obstacle segmentation individually adjusted
- Increase tracking stability
  - Must balance accuracy with both computation time *and* reactivity
  - Objects very close to the USV will be accelerating in the frame, not leaving much time to calculate moving averages

Further work

- Increasing accuracy and speed
  - Image preprocessing and prefiltering
  - Comparison with Kalman filter/statistical methods
- Combine imaging techniques
  - Correlated hybrid thermal/optical system
  - Different strengths (horizon, object segmentation)
- Integration with path planning
  - Field testing

Questions?

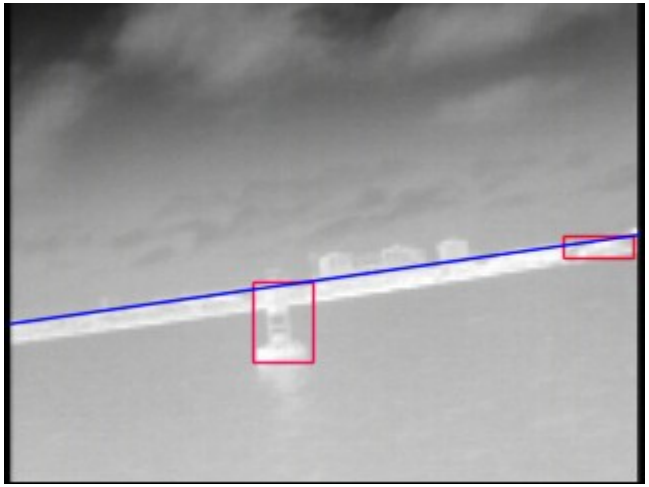

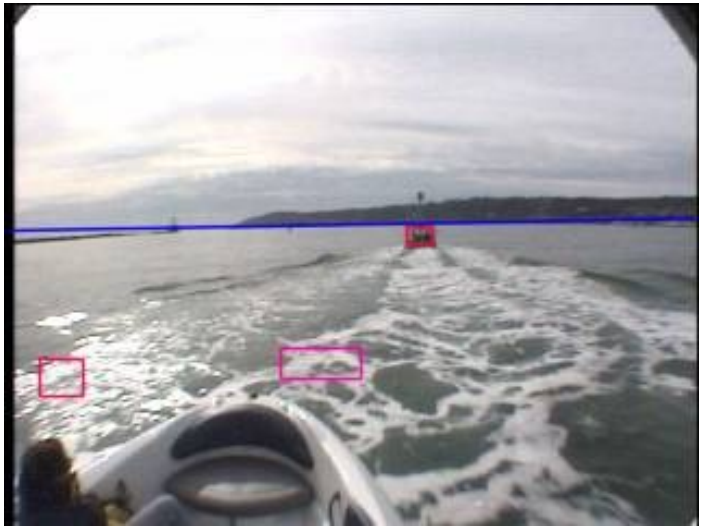

Daniel Donavanik

School of Engineering and Applied Science

University of Pennsylvania

ddonavan@cis.upenn.edu